Astronomy & Astrophysics manuscript no. Doppler
May 20, 2014
©ESO 2014

A multi-method approach to radial-velocity measurement for single-object spectra

M. David[1], R. Blomme[2], Y. Frémat[2], Y. Damerdji[3,4], C. Delle Luche[5], E. Gosset[4], D. Katz[5], and Y. Viala[5]

[1] Departement Wiskunde en Informatica, Universiteit Antwerpen; Middelheimlaan 1, B-2020 Antwerpen, Belgium
[2] Royal Observatory of Belgium; Ringlaan 3, B-1180 Brussels, Belgium
[3] Centre de Recherche en astronomie, astrophysique et géophysique. Route de l'Observatoire BP 63 Bouzareah - DZ-16340 Algiers, Algeria
[4] Institut d'Astrophysique et de Géophysique, Université de Liège; Allée du 6 Août, 17, B-4000 Liège, Belgium
[5] GEPI, Observatoire de Paris, CNRS, Université Paris Diderot; Place Jules Janssen 92195 Meudon Cedex, France

May 20, 2014

arXiv:1405.4303v1 [astro-ph.IM] 16 May 2014

ABSTRACT

Context. The derivation of radial velocities from large numbers of spectra that typically result from survey work, requires automation. However, except for the classical cases of slowly rotating late-type spectra, existing methods of measuring Doppler shifts require fine-tuning to avoid a loss of accuracy due to the idiosyncrasies of individual spectra. The radial velocity spectrometer (RVS) on the Gaia mission, which will start operating very soon, prompted a new attempt at creating a measurement pipeline to handle a wide variety of spectral types.
Aims. The present paper describes the theoretical background on which this software is based. However, apart from the assumption that only synthetic templates are used, we do not rely on any of the characteristics of this instrument, so our results should be relevant for most telescope-detector combinations.
Methods. We propose an approach based on the simultaneous use of several alternative measurement methods, each having its own merits and drawbacks, and conveying the spectral information in a different way, leading to different values for the measurement. A comparison or a combination of the various results either leads to a "best estimate" or indicates to the user that the observed spectrum is problematic and should be analysed manually.
Results. We selected three methods and analysed the relationships and differences between them from a unified point of view; with each method an appropriate estimator for the individual random error is chosen. We also develop a procedure for tackling the problem of template mismatch in a systematic way. Furthermore, we propose several tests for studying and comparing the performance of the various methods as a function of the atmospheric parameters of the observed objects. Finally, we describe a procedure for obtaining a knowledge-based combination of the various Doppler-shift measurements.

Key words. Methods: data analysis – Techniques: radial velocities – Stars: kinematics and dynamics – Surveys

1. Introduction

The importance of acquiring the radial velocity (RV) of extra-solar and extra-galactic objects is well known. Many techniques that do so, mostly based on a measurement of the Doppler shift of their spectra (but see e.g. Dravins et al. 1999, for an alternative), have been established during the past five decades, and detailed improvements are still being made. In the present era of large spectroscopic surveys such as RAVE (e.g. Steinmetz et al. 2006, 2009) and Gaia (e.g. Katz et al. 2004; Wilkinson et al. 2005), the ever-growing amount of available data requires increased automation in the scientific analysis, but this may entail diminished accuracy or reliability of the final results. The present paper revisits the question of Doppler shift measurement with the aim of reducing this risk. It provides the theoretical background of the way in which RV measurement will be done in the single transit analysis (STA) of the Gaia[1] radial velocity spectrometer (RVS) data, but its applicability is not limited to a particular telescope-instrument combination. The specific STA implementation will be discussed in a forthcoming paper.

1.1. Historical background

The idea of using the Doppler effect for determining stellar radial velocities dates back to nearly 150 years ago (Klinkerfues 1866; Sohncke 1867). The first successful attempt to bring it into practice was made by Huggins (1868). By the end of the 19$^{\text{th}}$ century the measurement of Doppler shifts was well-established (Cornu 1890) although there were still many technical issues to be solved, such as the stability of the spectrograph (Deslandres 1898). Until the 1960's one used the position of individual spectral lines, as is occasionally still being done for specific reasons (e.g. Andersen & Nordstrom 1983; Fekel 1999). Fellgett (1955) proposed a cross-correlation technique, already well established for analysing radar signals (e.g. Woodward & Davies 1950; Woodward 1953), which would enable the use of all the information contained in a spectrum, rather than the position of just a few lines. His ideas were developed in

[1] http://www.cosmos.esa.int/web/gaia

practice by Griffin (1967) and later perfected in dedicated instruments (Griffin & Gunn 1974; Baranne et al. 1979), all based on direct comparison of the spectrum with an optical mask.

In contrast to these analogue methods, Simkin (1974) proposed a digital cross-correlation; in spite of some initial shortcomings, it promised a huge advantage in flexibility, error control, and in particular efficiency, owing to the use of the fast Fourier transform (FFT) for calculating the cross-correlation function (CCF). This proposal was followed up by several reformulations or rederivations to adapt the generic technique for application to data sets with specific characteristics (e.g. Da Costa et al. 1977; Groote & Hunger 1977; Sargent et al. 1977). Later Tonry & Davis (1979) provided an alternative justification based on a χ^2 minimization, thus removing some apparent limitations of the method. Cross-correlation using FFT eventually became the main "off-the-shelf" method for RV measurement in many data reduction packages (e.g. Hill 1982; Kurtz & Mink 1998), reaching precisions better than $100\,\mathrm{m\,s^{-1}}$ for G- and K-type stars. We refer to it hereafter as the standard method. Its last significant reformulation is from Zucker (2003), who derives it from a maximum likelihood argument that enables him to provide a consistent estimate for the actual random error in a given measurement. He also shows how to combine the information from different orders of a spectrum without merging them. Although it was originally conceived for galaxies and single stars, the standard method has also been applied to spectroscopic binaries (e.g. Hill 1993; Ramm 2004). A formal extension to deal with double or triple spectra was elaborated by Mazeh & Zucker (1992), Zucker & Mazeh (1994) and Mazeh et al. (1995), and for quadruple spectra by Torres et al. (2007).

While the concept and the implementation of the standard method are fairly straightforward, its proper use in practice is less so. In fact it requires considerable fine-tuning to yield the best possible information, depending on the nature of the spectra involved. Late-type single-star spectra are not very demanding in this respect, but for early types fine-tuning may have to be different from one temperature class to the next in order to reduce spectrum mismatch errors, especially if the object is rotating fast (e.g. Verschueren et al. 1999; Griffin et al. 2000). Of course the availability nowadays of high-quality synthetic spectra to serve as templates alleviates this problem but even so, spectrum mismatch cannot be entirely excluded and may, in the absence of a large number of sharp lines, cause significant errors on the measured Doppler shift.

Several (likewise digital) alternatives for the standard method have been proposed; they could be classified roughly as minimum-distance (e.g. Weiss et al. 1978; De Loore et al. 1987; Zwitter et al. 2005), mask (Furenlid & Furenlid 1990; Baranne et al. 1996), phase-shift (Chelli 2000) or Pearson-correlation-coefficient (Royer 1999; Zverko et al. 2007) methods. Furenlid & Furenlid (1990) showed that the minimization of the sum-of-squares distance proposed by Weiss et al. (1978) is actually equivalent to the standard method, but the other alternatives are truly different (see also Sect. 2.4).

Recent developments were (and are being) driven by the quest for exoplanets and by asteroseismology, both of which were envisaged already in the 1990's but really took off around the turn of the millennium. These studies require measurement precisions better than $10\,\mathrm{m\,s^{-1}}$, which is close to the photon-noise limited precision (Butler et al. 1996; Bouchy et al. 2001) even for high-resolution spectra of bright slowly rotating late-type stars. Accuracy is not an issue here because one needs to detect only a variation of the radial velocity, which led Connes (1985) to coin the term accelerometry for this kind of work. Dedicated instruments were built with a focus on stability and on precise wavelength calibration using Th-Ar comparison spectra (Baranne et al. 1996; Lovis & Pepe 2007), absorption cells (e.g. Butler et al. 1996; Kambe et al. 2008, and references therein), and (most recently) laser frequency combs (e.g. Murphy et al. 2007; Steinmetz et al. 2008; Li et al. 2008; Cahoy et al. 2010), which could improve the precision of wavelength calibration to the order of $1\,\mathrm{cm\,s^{-1}}$.

The exoplanet search has also shifted observers' attention towards less massive stars (because their reflex motion must have a larger amplitude), which are best observed in the infrared (e.g. Seifahrt & Kaeufl 2008; Mahadevan & Ge 2009; Reiners et al. 2010; Valdivielso et al. 2010). Besides "plain" spectroscopy, new techniques based on a combination of interferometry and spectroscopy have emerged (e.g. Behr et al. 2009; van Eyken et al. 2010, and references therein). In most cases basically the standard method is still being used for the actual Doppler-shift measurement.

1.2. Measurement algorithms

The focus of the present paper is not, however, on such ultra-high precision work but on an approach that can yield reasonable (if not truly optimal) results for a wide range of spectral types without any object-dependent tuning so that it can be fully automated. This approach involves building an environment in which several methods are applied in parallel and the final outcome is a knowledge-based combination of their various results, the "knowledge" consisting of performance-test results obtained from simulated observations and stored in an auxiliary data base.

The fundamental assumption of any Doppler-shift measurement is that, for a given observed spectrum f one has a collection of template spectra t that are identical apart from a known Doppler shift and that differ from f only by the object's RV and by a random (observational) error on the flux samples. Furthermore one assumes that there is some measure of similarity or dissimilarity between the spectra, to be maximized, resp. minimized over the collection of templates. If no specific choice is intended we shall commonly call this measure a C-function. In general none of the available templates has exactly the same Doppler shift as the object, so an interpolation is required to obtain this velocity from the C-function samples; this interpolation is often termed centroiding. Altogether an algorithm for Doppler-shift measurement thus essentially consists of a (dis)similarity measure and a method of centroiding.

The main motivation for implementing more than one algorithm is as follows. All C-functions by definition reach their extremum at the same template velocity if the latter can be varied continuously and if object and template then match exactly (no noise, no spectrum mismatch etc.). However, they all respond differently to any deviation from this ideal situation, i.e. they differ in their sensitivity to the inevitable error sources (both random and systematic) and, if one has to deal with a large variety of spectra (as e.g. in the case of Gaia's STA) one cannot expect any sin-

gle method to be optimal for all of these. Furthermore, any significant difference (i.e. in excess of the expected random error) between the results they produce, may provide a useful indication of special care being required for the object at hand. Another possible advantage of using alternatives to the standard method, is that the latter requires rebinning the spectra to equidistant points in ln(wavelength), which may be undesirable if one wishes to exploit accurate knowledge of the random error distribution for individual observed flux samples (as, again, in the case of Gaia).

In Sect. 2 we select three C-functions and we review their foundation in a way that highlights the relations and differences between them. Sect. 3 discusses centroiding and the so-called model mismatch error it may entail. Sect. 4 deals with the problem of template mismatch. For each of the measurement algorithms an appropriate random-error estimate is chosen in Sect. 5. In Sect. 6 we discuss several tests for assessing the performance of the algorithms, as well as the storage of the test results. In Sect. 7 we use the latter to combine the various measurements resulting from the chosen methods. Finally, Sect. 8 offers some conclusions.

Note that we do not consider any corrections for the gravitational redshift, convective motions etc., so the term radial velocity used here actually refers to the "barycentric radial velocity measure" defined by Lindegren & Dravins (2003), not to the true line-of-sight velocity as determined from e.g. astrometric observations (Dravins et al. 1999; Lindegren et al. 2000). Such corrections, if required scientifically, will be assumed to be built into the library of synthetic templates one uses or to be added outside the framework of the Doppler-shift measurement.

2. Foundation of the C-functions

The C-functions we consider for a multi-method approach are the standard CCF (Sect. 2.3.1) because its computation using FFT is much faster than the others, the Pearson correlation function proposed by Royer (1999) (Sect. 2.3.2) because it is very flexible, not requiring the observed fluxes to match the template in an absolute sense, and a function derived from the maximum likelihood principle (Sect. 2.3.3). We chose to concentrate on algorithms that are applicable to as wide a variety of spectra as possible, so we do not consider a phase-shift method because its applicability would be limited to very high S/N, nor a digital-mask method because it would require high resolution. Nevertheless, these algorithms do have their merits and, if it is judged useful for some reason, they could easily be included in a multi-method approach.

2.1. Notations and definitions

Throughout this paper we use λ to represent wavelength in general, λ_n for the central wavelength of the $n^{\rm th}$ bin of an observational grid, Δ_n for half the width of this bin (in general the width will vary along the spectrum), $f(\lambda)$ for the spectrum emitted by an observed object, f_n for its $n^{\rm th}$ flux sample (which is of course the only real information we have about this spectrum), σ_n for the random-error estimate on this sample, $t(\lambda; v)$ for a template spectrum (see Section 2.2.2) that has been calculated for a source moving with radial velocity v (called the template velocity), and $t_n(v)$ for the $n^{\rm th}$ sample of the latter, binned to the observed wavelength grid, i.e.

$$t_n(v) = \int_{\lambda_n-\Delta_n}^{\lambda_n+\Delta_n} t(\lambda; v)\mathrm{d}\lambda \ . \tag{1}$$

Then, assuming that the template exactly matches the source at hand, we can state that the observed fluxes are related to the template by

$$f_n = a_{\rm s} t_n(v_{\rm s}) + d_n \ , \tag{2}$$

where $a_{\rm s}$ is merely a scale factor (e.g. for the conversion of units or to allow for some uncertainty in the continuum level of the object) and $v_{\rm s}$ is the radial velocity of the source, both to be determined from the data, while d_n is a random variable with zero mean, representing the noise in the $n^{\rm th}$ observed sample. If the scale factor $a_{\rm s}$ is known a priori one can include it in the definition of the template and put $a_{\rm s} = 1$ in the expressions to be derived below.

The term data segment may refer to a set of observed fluxes $(f_1, f_2, \ldots, f_N)$ for a given object as well as to a set of "predicted" fluxes $(t_1(v), t_2(v), \ldots, t_N(v))$.

If a data segment has been rectified so that its (pseudo) continuum level is constant, and if its ends lie within a line-free region of the spectrum (see Verschueren & David 1999, for more details), it will be termed regular. Even though in practice we cannot expect all spectra to be truly regular, this notion will be useful in the discussion below.

2.2. Properties of the input data

2.2.1. Observed flux

We assume that

1. the observed spectra have been prepared (subtraction of bias, dark current, and celestial background; detection of cosmic-ray effects and CCD blemishes; extraction; wavelength calibration) in both normalized and non-normalized form, the latter implying that the value of each sample is expressed as a number of photons (but note that this number is generally non-integer after the data preparation);
2. the random errors on the flux samples are statistically independent;
3. the probability density function (PDF) of this random error is Gaussian with known "real" dispersion for each sample, obtained e.g. from standard error propagation taking account of all relevant noise sources (photon count, read-out, bias, dark current, background etc.); however, to provide for a situation where such error estimates are not available we also discuss two alternative models for this PDF (see Sect. 2.3.3);
4. flux samples affected by cosmic-ray hits or CCD blemishes have been given some reasonable value but they are also marked so that (where necessary) they can be excluded from the calculations;
5. any wavelength regions containing spectral features that originate from matter that does not share the motion of the source (e.g. interstellar lines) have been properly identified and the flux samples within them are marked as above;
6. the spectra are not noise-filtered because that would inevitably destroy the assumed statistical independence of the flux samples.

We do not presume that the spectra are contiguous: one may determine separate C-functions for e.g. the orders of an echelle spectrogram, different CCDs (as in the case of Gaia) or disjoint wavelength regions (e.g. to eliminate interstellar or atmospheric spectral features); afterwards these C-functions can be averaged (Zucker 2003) or, for those of Sect. 2.3.3, summed.

2.2.2. Template

In principle we assume that

1. the atmospheric parameters (APs) of the observed object are known and a corresponding synthetic spectrum $s(\lambda)$ is available with virtually infinite resolution; the former condition is seldom realistic so in Sect. 4 we discuss the consequence of relaxing it;
2. the projected rotational velocity of the object is known and the synthetic spectrum has been broadened accordingly;
3. the template spectrum $t(\lambda, v)$ is obtained by applying the Doppler shift, smooth[2] interstellar extinction and all known instrumental effects (optical distortion, line-spread function (LSF), CCD recording, ...) to the (rotationally broadened) spectrum $s(\lambda)$; it is provided with and without normalization;
4. the template is calculated on a wavelength range that extends sufficiently far beyond the observed one so that after any foreseeable Doppler shift, the former still completely covers the latter.

The second assumption is made for the following reason. Tonry & Davis (1979) have shown how line-broadening due to velocity dispersion in the spectrum of a galaxy may be determined along with the radial velocity in a two-dimensional optimization. In principle their approach can be applied also to rotational broadening in a stellar spectrum (e.g. Díaz et al. 2011). However, this may entail some correlation between the two measured values, which is often undesirable from the point of view of error propagation. Therefore we assume that $v \sin i$ is obtained from an independent measurement, as e.g. in Royer et al. (2002); Gray (2005); Chen et al. (2011). The third assumption ensures that Eq. (2) is strictly applicable. This may not be required by some of the methods we discuss here, but of course the template must be the same for all of them if we want a valid comparison between their respective outcomes.

Formally the template may be represented as

$$t(\lambda; v) = \int_0^\infty G(\lambda, \lambda_s(1 + \frac{v}{c}))s(\lambda_s)\mathrm{d}\lambda_s \ , \tag{3}$$

where G is a Green's function, λ is the wavelength in the observer's rest frame and λ_s is the wavelength in the rest frame of the synthetic spectrum.

This form implies that the integration should be performed for every value of v considered. However, in many cases this integral is well approximated by a simple convolution of s with the LSF, multiplied by a scale factor. If the LSF is sufficiently narrow one can show that actually

$$t(\lambda; v) \approx \frac{1}{1 + \frac{v}{c}} t(\frac{\lambda}{1 + \frac{v}{c}}; 0) \tag{4}$$

so that the convolution must be computed only once and Eq. (1) can be written as

$$t_n(v) = \int_{\frac{\lambda_n - \Delta_n}{1+\frac{v}{c}}}^{\frac{\lambda_n + \Delta_n}{1+\frac{v}{c}}} t(\lambda; 0)\mathrm{d}\lambda \ . \tag{5}$$

In situations where Eq. (4) is no longer valid for large template velocities, e.g. due to charge distortion effects caused by radiation damage (which cannot be accounted for in a convolution), one can compute Eq. (3) on a coarse velocity grid and apply an approximation such as Eq. (4) only to deviations from these grid values.

2.3. C-functions

2.3.1. Cross-correlation – the standard method

The cross-correlation function (CCF) is a descriptive tool developed for investigating stochastic processes. In her pioneering paper on the digital measurement of Doppler shifts, Simkin (1974) referred to the standard literature in this field (e.g. Jenkins & Watts 1969; Bendat & Piersol 1971) to introduce the use of cross-correlation for Doppler-shift measurements, stressing the necessity of rectifying (Jenkins & Watts 1969) the spectra and of rebinning them to a grid with constant step in $\ln \lambda$ (see Appendix A) so that they can be considered as a so-called stationary random process.

In our notation, taking account of Eq. (A.4), the CCF as defined by Tonry & Davis (1979) is

$$C_{\mathrm{cc}}(m) = \frac{\sum_{n=1}^{N}(f_n - \bar{f})(t_{n-m}(0) - \bar{t}(0))}{\sqrt{\sum(f_n - \bar{f})^2 \sum(t_{n-m}(0) - \bar{t}(0))^2}} \ , \tag{6}$$

where $\bar{}$ indicates the average over a data segment. Note that f_n and t_n are considered here as periodical functions of n with period N; this implies that indices $n - m < 1$ or $n - m > N$ correspond to wrapped-around values of t_n with $1 \le n \le N$. This wrap-around may cause a systematic error (Simkin 1974) unless both data segments are regular and the shift does not exceed the length of the line-free region at either end. In case of larger shifts one may therefore need to make a crude measurement first and then use an adapted data segment for the template to do the final measurement.

Although we adopt this definition, for stellar spectra we do not follow most of the data preparation that Tonry & Davis (1979) advocate:

- instead of continuum subtraction we prefer normalization, which serves the same purpose while conserving the relative strength of the lines;
- both high- and low-frequency filtering are omitted;
- we prefer to avoid endmasking (or apodization) because this introduces a structure in the observed spectrum that does not share the Doppler shift of the stellar features, thereby possibly causing a systematic error (Verschueren 1991, p. 131); the purpose of endmasking is better served by adapting the wavelength range (if possible) to make the data segments regular.

In the standard method one does not evaluate the sum in Eq. (6) directly, but via its Fourier transform, using the so-called Fast-Fourier technique; in that case the number of flux samples in each data segment must be a power of two.

[2] i.e. insofar as it can be described by a factor that is constant or at most slowly varying over the wavelength range considered.

This may be achieved by adding, at either end, a number of samples having the continuum value (an operation also known as "zero-padding") or by adapting the stepsize in $\ln\lambda$, provided that the adopted stepsize can be kept close to the average observed bin width.

If the observed spectrum contains samples marked as invalid (see Sect. 2.2.1), one has two options: either to use some reasonable replacement value for the invalid samples or to split up the data segment into smaller ones that do not contain any "bad" data. The latter option is preferable if several adjacent pixels are affected but it is hardly practicable in an automated treatment.

2.3.2. Pearson correlation

Royer (1999), while referring to Simkin (1974) and Tonry & Davis (1979), actually adopted a quite different approach using the classical Pearson correlation coefficient as a measure of the similarity of the data segments representing respectively the object and the template calculated for a particular template velocity v. This implies that he considered these data segments merely as ordered sets of measurements of two correlated random variables or, rearranged in the form $\{(f_n, t_n(v))|n = 1, 2, \ldots N\}$, as an ordered set of values drawn from a bivariate population with a certain correlation. The distribution of this population would be nondescript[3] but we note that anyway all its moments would be formally finite. If the observed spectrum is also noiseless and different only by the radial velocity of its source, the correlation will depend only on the object–template velocity difference and on the wavelength binning, becoming a strict equality if the velocities are equal.

Royer (1999) prepared his data with the logarithmic sampling Eq. (A.1), but this is not required in principle. In fact we can define the resulting C-function in general as

$$C_{\rm pc}(v) = \frac{\sum_{n=1}^{N}(f_n - \bar{f})(t_n(v) - \bar{t}(v))}{\sqrt{\sum_{n=1}^{N}(f_n - \bar{f})^2 \sum_{n=1}^{N}(t_n(v) - \bar{t}(v))^2}} \ , \qquad (7)$$

This expression presumes nothing about the wavelength bins in Eq. (1), except that the grid $(\lambda_1, \lambda_2, \ldots, \lambda_N)$ is common to both spectra. Sample indices n that have been marked as invalid in the observed spectrum (see Sect. 2.2.1) are simply to be omitted from all sums (also from those containing only template data). Hereafter Eq. (7) will be termed the Pearson correlation function (PCF) as distinct from the CCF.

Unlike the CCF, the PCF does not require that the data segments are regular or even rectified[4] so it should be applicable also where e.g. rectification is problematic. Nevertheless one should be aware that its sensitivity is likely to be degraded if the ratio of the continuum levels of object and template varies considerably over the spectral range.

We also note that, in its general form with an arbitrary wavelength grid for the observed data, the PCF can be evaluated only in velocity space; this is in fact the case we consider hereafter.

[3] Incidentally, it would look Gaussian-like in most parts of late-type spectra, but otherwise its shape could be very different.

[4] This does not imply that, for a given observed spectrum, the Doppler-shift value measured on the normalized form should be strictly identical to the one obtained with the non-normalized form, but the difference is usually much smaller than the common measurement errors.

2.3.3. Maximum likelihood / minimum distance

The noise in different samples is uncorrelated but of similar origin so we can consider a data segment as a measurement of a set of N independent random variables that are distributed similarly, their frequency functions differing only by their expectation e_n and dispersion σ_n where, according to Eq. (2),

$$e_n = a_{\rm s} t_n(v_{\rm s}) \quad , \quad \sigma_n = {\rm stdev}(d_n) \ . \qquad (8)$$

We note ${\rm pd}(f_n|a, v)$ for their individual PDF. Since the form of this function is known or can be reasonably hypothesized, it is justified to use the maximum likelihood (ML) principle for estimating the parameters $a_{\rm s}$ and $v_{\rm s}$ (Zucker 2003). Given a data segment with N samples, the likelihood function is given by

$$L(a, v|f) = \prod_{n=1}^{N} {\rm pd}(f_n|a, v) \ . \qquad (9)$$

Invalid sample indices should be omitted from this product (and thus from all sums derived from it hereafter). The parameter values $(\hat{a}, \hat{v})$ that maximize this likelihood constitute a consistent asymptotically minimum-variance estimator for the set $(a_{\rm s}, v_{\rm s})$ (see e.g. Martin 1971, and references therein).

Whereas the PCF essentially measures the similarity of the flux gradients, the likelihood function measures the similarity of the fluxes themselves. Therefore a method based on the ML principle is also a minimum distance method, which is a more significant name in the present context.

Zucker (2003) assumes exclusively that the distribution is Gaussian and that the dispersions σ_n are all equal but unknown; he determines this unknown value $\sigma_{\rm s}$ from the data along with the other parameters. However, in Sect. 2.2.1 we made other assumptions so we need to derive the appropriate equations for our case. If the PDF is Gaussian with known dispersion σ_n for each sample, then

$${\rm pd}(f_n|a, v) = \frac{1}{\sqrt{2\pi}\sigma_n} e^{-\frac{(f_n - a t_n(v))^2}{2\sigma_n^2}} \ . \qquad (10)$$

Taking the logarithm of Eq. (9) one easily sees that maximizing the likelihood is equivalent to minimizing

$$C_{\rm md}(a, v) = \sum_{n=1}^{N} \frac{(f_n - a t_n(v))^2}{\sigma_n^2} \qquad (11)$$

which, incidentally, equals the classical χ^2 statistic for a goodness-of-fit test. Obviously this expression can in fact be seen as an error-weighted squared "distance" between object and template.

As pointed out before, if a is known independently from the minimization problem, it may be omitted from this expression. Otherwise the two-dimensional minimization to determine $(\hat{a}, \hat{v})$, can be reduced to one dimension by using the fact that the derivative $\frac{\partial C_{\rm md}}{\partial a}$ must vanish at the minimum; this leads to

$$\hat{a} = \frac{\sum \frac{f_n t_n(\hat{v})}{\sigma_n^2}}{\sum \frac{t_n^2(\hat{v})}{\sigma_n^2}} \ . \qquad (12)$$

Substituting this in Eq. (11) we find that $\hat{v}$ is the value that maximizes

$$C_{\rm md1}(v) = \frac{\left|\sum \frac{f_n t_n(v)}{\sigma_n^2}\right|}{\sqrt{\sum \frac{t_n^2(v)}{\sigma_n^2}}} \ . \tag{13}$$

Having determined $\hat{v}$, one finds $\hat{a}$ from Eq. (12).

If, for some reason, the error estimates σ_n are not available while, on the other hand, it would not be justified to assume that they are all equal, one can adopt an alternative model for the flux distributions, such as

- a Gaussian with a predicted dispersion describing photon noise (derived from the template) and possibly a stationary additive error source (e.g. read-out noise);
- the Poisson probability function, applicable if one can assume that the recorded flux samples represent a pure photon count.

Appropriate minimum-distance equations for these models are given in Appendix B.

2.4. Similarity between different C-functions

While all the functions defined above are obviously different from each other, there is some (at least superficial) similarity between them, and it is worthwhile to explore this in more detail. If the binning is logarithmic, if $t_n(v)$ is calculated for the discrete velocities given by Eq. (A.3) and if both the object and the template (for all template velocities considered) are regular as defined in Sect. 2.1, then $\sum t_n(v)$ and $\sum t_n^2(v)$ are constant over a limited velocity range around the object's velocity; in that case neither the denominator in Eq. (7), nor $\bar{t}(v)$ in the numerator, influence the extremum position so that $C_{\rm pc}(v)$ becomes equivalent to $C_{\rm cc}(v)$ (i.e. both must yield exactly the same value for $\hat{v}$). Likewise for such a dataset, if moreover all error estimates are equal (i.e. $\sigma_n = \sigma_0 \ \forall\, n$) then also $C_{\rm md1}(v)$ becomes equivalent to the former two.

With any new implementation of these (and possibly other) methods it is advisable to test the above, first with noiseless data so that the results must agree within the known bounds of machine precision, then with some noise added to obtain a first indication of how each method reacts to that. If these tests are satisfactory one can consider gradually more realistic cases and so acquire a good understanding of the differences between the measurements and of their relation to the characteristics of the observed spectra (atmospheric parameters, rotational broadening, resolution, S/N-level, ...).

3. Measurement

3.1. Identifying the peak

In many applications the range of velocities one expects to encounter will be very wide; nevertheless on physical grounds it must be possible to estimate reasonable boundaries for this range, say, $[v_{\rm min}, v_{\rm max}]$. Then a first (coarse) estimate of the Doppler shift can be made as follows.

With the methods operating in velocity space one calculates the C-function over the whole velocity range $[v_{\rm min}, v_{\rm max}]$ on a relatively coarse grid with step size sufficiently small to detect all local extrema (e.g. $10\,{\rm km\,s^{-1}}$ for spectra with instrumental resolution 11500) and one locates the highest/lowest value. Then, reducing the range to one coarse step to either side of the latter and the step size with some factor (e.g. 10), the location is improved. If necessary a further refinement may be obtained in the same way. The actual measurement is done by centroiding the extremum within the resulting (small) set of C-function samples.

On the other hand, the standard method implies that a fixed segment of the template spectrum is "shifted" along the observed one. As a consequence, if the source has very large RV, the template segment cannot match the observed spectrum very well (even at the correct shift) because either will contain a part of the spectrum that the other does not. Moreover the CCF is not truly refined by choosing a smaller velocity step (i.e. a smaller bin width in $\ln\lambda$) because that would require resampling also the observed spectrum to smaller bins, which may imply an interpolation of mere noise fluctuations. For these reasons we proceed as follows.

First consider the two "extreme" template data-segments, obtained with $v_{\rm min}$ and $v_{\rm max}$ resp. as the template velocity. Cross-correlating each of these with the observed spectrum should yield
- either two shift values that are comparable
- or widely different values, one of which corresponds to a much more significant peak in the CCF than the other (e.g. difference in height exceeds five times the expected random error on the CCF values).

If (in an automated treatment) neither of these situations occur then manual intervention is advisable. With the above prescription one can obtain an indication of whether the RV is much smaller than the width of the whole interval, or comparable to it; in the latter case it is best to use a template data segment calculated on a $\ln\lambda$ interval that has been pre-shifted according to the coarse estimate, so that the residual shift to be found by centroiding is relatively small.

It should be noted here that with very noisy spectra the highest/lowest C-function value identified as "the peak" may be totally unrelated to the Doppler shift so that the above procedures (and indeed the measurement as such) become meaningless. This situation, where the observed extremum position may deviate much more from the true Doppler shift than predicted by any error estimate, is discussed in Sect. 5.2.

3.2. Centroiding

Centroiding is essentially the interpolation between a number of discretely sampled C-function values with a view to find the position of the extremum of the underlying (continuous) function. Ramm (2004, section 5.7.2) and Kurtz & Mink (1998) list a number of interpolating functions that can be used for this purpose. However, any interpolating function that does not equal the underlying function (up to a shift) may itself be a source of error, the so-called model-mismatch error (MME, David & Verschueren 1995), and generally makes the result sensitive to the number of samples (or "fit points") used for determining the parameters of the interpolating function.

As in any interpolation problem the accuracy of the result improves if the interpolating function better resembles the underlying function. Therefore any information one

has on the intrinsic nature of the C-function, may be useful in choosing an appropriate model. For instance if one knows that the C-function is intrinsically symmetric with respect to its extremum, obviously the interpolating function should have the same property. And if moreover one has reasons to expect that the C-function is close to e.g. Gaussian (as with slowly rotating G stars) one may try that as a model. Nevertheless it has been noted (Allende Prieto 2007) that even here the Gaussian (which in principle is more robust against noise) does not always perform better than a simple parabola. This is related to the fact that a Gaussian requires fit points down to the base level, which do not contain any information on the Doppler shift.

The number of fit points to be used, depends on several circumstantial factors. First of all, it is obvious that "clean" information on the Doppler shift resides only very close to the centre of the C-function peak or trough, and certainly not farther away from it than the HWHM of the sharpest lines in the observed spectrum.

If the C-function sampling step is much smaller than this maximum distance, one may be tempted to use as many fit points as possible, to suppress the effect of random errors, but this makes sense only if the random error on adjacent C-function values is statistically independent; that may be the case, approximately, in the standard method where the C-function sampling is tied to the observed wavelength bin, but it can hardly be true with the velocity-space methods, which refine the velocity step to a fraction of the value corresponding to the observed wavelength-bin width. Moreover, if the model does not match the intrinsic (i.e. noise-free) shape of the C-function, systematic deviations between those two will gain influence when fit points farther away from the extremum are used.

As a general rule, assuming we do not know the exact shape of the C-function, we thus find that it is best to use the simplest possible model (i.e. a parabola) and the minimal number of fit points (i.e. 3), unless there is clear evidence indicating otherwise; this conclusion fully agrees with the results of Allende Prieto (2007). The MME in this case, assuming the C-function is locally symmetric, is discussed in Appendix C.1.

One may object to the use of a parabola because this is a symmetric function while the velocity-space C-functions are slightly asymmetric so that we would knowingly cause a MME. However, the random error on the C-function samples generally introduces a much stronger asymmetry, so that the use of e.g. a third-degree polynomial would mean that the model will mainly pick up noise effects, degrading the actual position measurement (cf. Allende Prieto 2007). If there is a compelling reason to allow some asymmetry in the model, one should use at least a polynomial of even degree (Verschueren & David 1999).

On the other hand, knowing that the use of a parabola for centroiding an asymmetric C-function may thus cause a somewhat larger MME than the basic one discussed in Appendix C.1, we need a way to check whether it is in fact still negligible in a given situation. Therefore in Appendix C.2 we derive (under fairly mild conditions) an upper bound for this error.

4. Template mismatch error

4.1. General remarks

The template used for measuring the RV of a given single star, is derived from a synthetic spectrum, selected from a library using the object's atmospheric parameters and rotationally broadened with an estimate of the object's $v \sin i$.

Under most circumstances, even if some of the APs could be determined with very high accuracy, their value will not correspond exactly to any of the available synthetic spectra. Given the almost inevitable lack of knowledge on several other factors determining the details of an observed spectrum, it seems pointless to consider interpolating between library spectra. Therefore we assume that simply the one nearest (in AP space) to the observed one, will be used. This means that in principle each of the chosen template's APs could be in error by 1/2 its step size in the library, even under otherwise ideal circumstances.

Almost any mismatch between template and object may bias the RV measurement. Such bias is usually referred to as the template mismatch error (TME). In principle this is a systematic error, since it will have the same sign and order of magnitude for all observed objects with the same intrinsic AP values. Nevertheless, the uncertainty on the APs being random, the user of the RV measurements may decide to treat the TME as if it were an additional random error.

To estimate the TME, North et al. (2002) measure each object against a set of templates rather than just the best matching one; the resulting set of RV-values will then provide an estimate of the TME. This would surely suffice if one has to deal with only a small number of objects, but e.g. in survey work one may prefer a quick estimate obtained from a table or from a simple model, even though inevitably any such estimate would be less realistic.

Therefore in this section we propose a strategy for studying, with a given synthetic library and in a given region of AP space, how the TME may behave both as a function of the AP errors (later referred to as a local model) and as a function of the spectral type of the object (a global model). This information should help a user of the RV measurements to assess the importance of the TME in view of the accuracy requirements and of the random uncertainty of the RV.

In the examples below we consider only the effect of an uncertainty in three APs and in $v \sin i$, but it is clear that other sources of mismatch, such as the convective deformation of line profiles (see e.g. Gullberg & Lindegren 2002, and references therein) could be treated in a similar way.

4.2. Modelling strategy

4.2.1. The AP grid

The dominant APs determining the spectral features are the temperature ($T_{\rm eff}$), gravity ($\log g$), and metallicity ([Fe/H]). Other APs such as [α/Fe], turbulent velocity, the abundance of specific elements or molecules, magnetic field strengths, etc. may be relevant in particular cases, depending on the required accuracy, but in the present global assessment they will be ignored. Thus we limit ourselves to a 3-dimensional AP-subspace within which the triplets ($T_{\rm eff}, \log g$, [Fe/H]) characterizing the available library spectra, define a grid with (generally) variable stepsize.

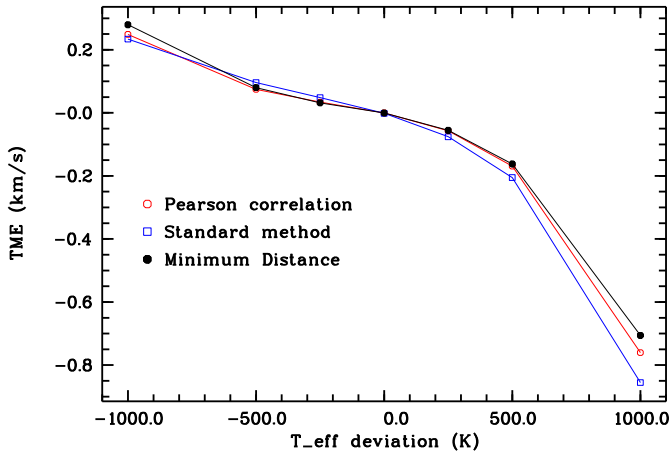

Fig. 1. Simulated TME for $c = (5500\,\mathrm{K}, 4\,\mathrm{dex}, 0\,\mathrm{dex}, 20\,\mathrm{km\,s^{-1}})$ with a local grid defined by Eq. (14) and $x_2 = x_3 = x_4 = 0$ while $x_1 = -4, -2, -1, 0, 1, 2, 4$.

Unlike the APs, the $v \sin i$-value to be applied is not technically limited to a discrete set; however, this value inevitably has some measurement error and we need to investigate how that may affect the final RV-error. Therefore, in the discussion below we shall treat $v \sin i$ on an equal footing with the atmospheric parameters, considering its estimated value as situated on a grid with local stepsize equal to twice the uncertainty of this value. As a consequence, formally we shall consider a 4-dimensional grid with each node characterized by a set of values $c = (T_{\rm eff}, \log g, [\mathrm{Fe/H}], v \sin i)$.

In physical terms this grid is not uniform, but we shall use it mainly for bookkeeping, so whatever the actual amount that is represented by a given edge, we consider the latter to have length "1" in the appropriate units. Assuming for the sake of simplicity that the AP steps are locally constant, we arrange these in a local "units" vector $u = (u_1, u_2, u_3, u_4)$, e.g.

$$u = (250\,\mathrm{K}, 0.5\,\mathrm{dex}, 0.5\,\mathrm{dex}, 10\,\mathrm{km\,s^{-1}}) \ . \qquad (14)$$

Noting (c_1, c_2, c_3, c_4) for the node chosen as the origin, any set of AP's in the latter's vicinity can be characterized by a vector $x = (x_1, x_2, x_3, x_4)$, $x_i \in \mathbb{R}$, the actual parameter values being given by $a_k = c_k + x_k u_k \ , \quad k = 1, \ldots 4$. With these definitions we can think of the grid as a simple hypercubic lattice.

4.2.2. Simulating the TME

We choose a given node in the AP grid as its origin and consider the corresponding template spectrum, prepared as in Sect. 2.2.2 at zero velocity, as an "observed" spectrum. Then we measure the latter's Doppler shift against the templates corresponding to the neighbouring nodes; all the spectra involved should be noiseless, since we are studying a systematic effect here. The resulting values constitute a local (discrete) sampling of the TME as a function of AP deviations $\Delta a_k = x_k u_k \ , \quad k = 1, \ldots 4$. Fig. 1 shows a simple example (only $T_{\rm eff}$ is variable) of such a function; notice that the effect of deviations with $|x_1| \leq 2$ could be described by a $2^{\rm nd}$-degree polynomial in x_1 but that larger deviations would require a $3^{\rm rd}$-degree polynomial. The following sections discuss the modelling of such functions.

4.3. Local model

Lacking any theoretical information about the functional dependence of the TME on the AP deviations, we can only assume that it is analytic so that the TME may be approximated by a truncated Taylor expansion. Approximations of increasing order can be thought of as taking into account "cross-talk" between different APs (i.e. the deviation in one parameter modifying the effect of a deviation in another parameter) of increasing complexity. E.g. in the first-order the effect of simultaneous deviations in several parameters is described simply as the sum of their individual effects; in a second-order approximation the combined effect of pairwise deviations is described more accurately, etc. We shall assume that a second-order approximation,

$$P^{(2)}(x) = \sum_{k=1}^{4} C_k x_k + \sum_{k=1}^{4} \sum_{l=1}^{k} C_{kl} x_k x_l \ , \qquad (15)$$

is sufficient for the purpose at hand. As one can see from Fig. 1, this assumption limits the size of the deviations for which the model is valid.

The coefficients C_k and C_{kl} in $P^{(2)}(x)$ can be obtained from TME measurements for the nodes characterized by a vector x that has three or two components equal to zero while the remaining ones are ± 1. There are 32 of these for the 14 coefficients, but we prefer to keep this redundancy and to obtain the coefficients from a least-squares fit to the larger set of data, because a matching number of nodes would be less symmetrically distributed over the hypercube and because in practice some of the nodes may be unavailable anyway owing to local incompleteness of the synthetic spectrum library. Actually such incompleteness may cause the system to be underdetermined even if nominally the number of data exceeds the number of unknowns; therefore in the polynomial fit we use a procedure based on the so-called Moore-Penrose pseudo-inverse (e.g. Penrose 1955, 1956), which allows the user to make an informed decision on which monomials to remove from the model if some coefficients cannot be uniquely determined.

Once the coefficients for the second-degree model have been obtained, the latter can be validated by verifying that, if applied to all adjacent nodes (of which there are 48 in addition to the 32 used for the fit), its residuals do not exceed a pre-set threshold. If the model is found to satisfy this requirement it may be trusted to predict the TME on an actual RV measurement in which the AP deviation is within the range described by the local units vector u.

4.4. Global model

The coefficients in the local model vary with the APs defining the central node. This variation should be investigated as well. Depending on the outcome of such a study, one may decide either to search for a global model or to establish a look-up table for the polynomial coefficients. A global model would mean that the 14 coefficients are represented, in their turn, by a function of the APs (c_1, c_2, c_3, c_4) of the central node. However, exploratory calculations (see e.g. Fig. 2) indicate that such a function could not be a low-order polynomial if it is to be valid over at least a sizeable region of AP space. Therefore we believe it is better to define a relatively coarse grid covering all of the relevant AP space, determine the coefficients in Eq. (15) for each of its

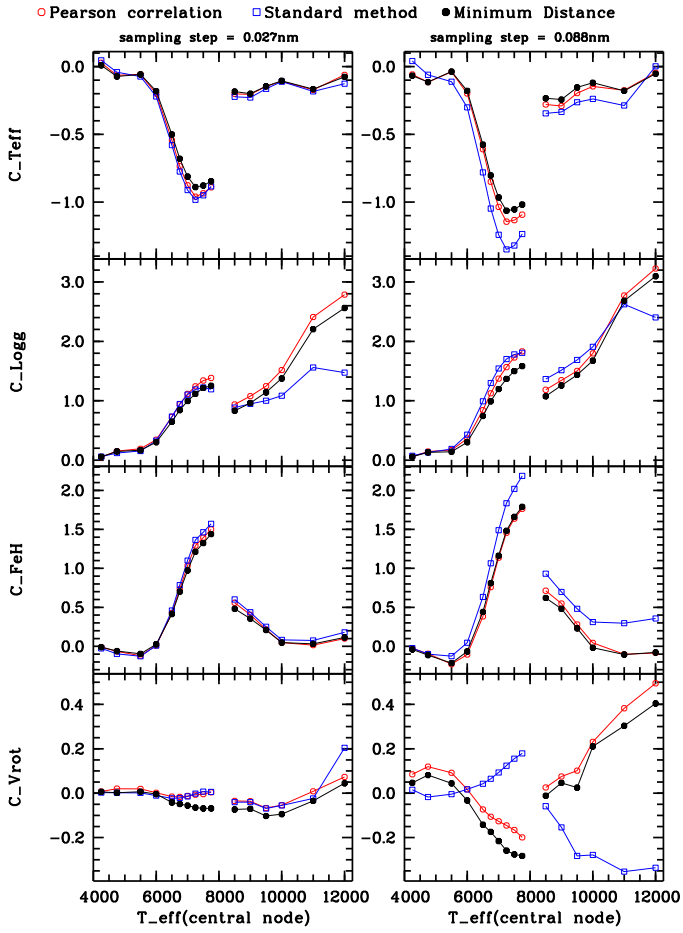

Fig. 2. The first-order coefficients in Eq. (15) for a set of central nodes with $\log g = 4$, [Fe/H]$= 0$, and $v \sin i = 20$, but with different $T_{\rm eff}$-values as indicated on the x-axis. The y-axis labels identify the deviating AP in each frame. The ordinate values represent the TME-contribution (in $\mathrm{km\,s^{-1}}$) of this AP if it deviates by one step (see Eq. (14)) from the central node. The run of the coefficients is given for slightly oversampled (left frame) and slightly undersampled (right frame) spectra with an instrumental resolution of 11500, and for the three measurement methods as indicated on top. The gap at 8000 K is due to the fact that different atmospheric models were used for temperatures above and below this value; incidentally it illustrates the sensitivity of Doppler-shift measurements to the templates used.

nodes and store these in an auxiliary database. When the local model is required for any other node, its coefficients may then be approximated by interpolation between those on the coarse grid. Obviously such a database will be valid only for a given telescope-instrument combination, i.e. for a particular mean instrumental resolution, wavelength range and –sampling step.

Incidentally, Fig. 2 also indicates that the TME may differ significantly between measurement methods, which is something one may wish to take into account in the combination of their respective results.

5. Random errors

5.1. Internal error estimators

Besides the various algorithms for Doppler-shift measurement, also several ways of estimating the random error on this measurement have been proposed; the term internal error is borrowed from Tonry & Davis (1979). Most of these were designed in accordance with a specific C-function and with certain assumptions about the data, so some care is due if one wishes to apply them in a different context. We now discuss their applicability under the particular requirement that they be realistic (at least in order of magnitude) as well as suitable for a wide range of spectral structures.

Tonry & Davis (1979) estimated the error on the radial velocity using a measure of the asymmetry of the CCF peak, based on the argument that this peak is intrinsically symmetrical if the two spectra are noiseless and match each other exactly. Their estimate thus describes not only the random error but also (at least part of) the systematic error caused by spectrum mismatch. However, Verschueren & David (1999) demonstrated that it cannot be valid for spectra that exhibit relatively broad features such as the Balmer lines in the visible region of early-type spectra, the CaII triplet or the Paschen lines in the far red etc.

Murdoch & Hearnshaw (1991) proposed an error estimate based on the standard expression for the error on the least-squares estimate applied in centroiding the C-function peak. However, in doing so they presumed that the random error on neighbouring C-function samples is uncorrelated whereas actually it is highly correlated in general (Jenkins & Watts 1969). In the case of the PCF and of the minimum-distance C-functions, the random error is in fact almost identical for adjacent C-function samples if these are calculated with a velocity step that is much smaller than the one corresponding to the step of the observed wavelength grid. Therefore this error estimate is not applicable here.

Brown (1990) derived a lower bound for the random error, assuming that there is only photon noise. However, in most applications a realistic error estimate, rather than a mere lower bound is required. We verified in a number of cases that Brown's estimator indeed yields values that are too small to be relevant in practice.

After extensive tests we found that the best error estimator to be used with both the CCF and the PCF, is the one proposed by Zucker (2003):

$$\sigma_v^2 = \frac{1 - C^2}{NCC''} \ , \tag{16}$$

where N is the number of wavelength bins while C and C'' represent respectively the functional value and the second derivative of the C-function ($C_{\rm cc}$ or $C_{\rm pc}$), evaluated at its centroid position.

For the minimum-distance functions we follow the approach of Cash (1979). Each of these C-functions equals (up to some constant terms which have been omitted) the statistic $-2 \ln L(a, v|f)$ and it can be shown that the difference ΔC defined as

$$\Delta C(v) = C(v) - C(\hat{v}) \ , \tag{17}$$

where C stands for any of the functions $C_{\rm md}(v)$, $C_{\rm mdp}(v)$ or $C_{\rm mdc}(v)$ (see Appendix B), is distributed approximately as χ^2 with one[5] degree of freedom. Thus we can define a

[5] Even in the case where a is also adjustable: in fact we are not interested in the joint confidence region for both parameters but only in the overall uncertainty on v (Press et al. 1986).

classical "one-σ" confidence region as the interval $[v_1, v_2]$ where $\Delta C(v_1) = \Delta C(v_2) = 1$. This interval needs not be symmetrical with respect to $\hat{v}$, so we adopt $\sigma_v = \max|v_i - \hat{v}|$ as the internal error estimate for a minimum-distance measurement.

5.2. Very faint objects

The internal error estimators all use (one way or another) the shape of the C-function in the immediate vicinity of its maximum. However, as the continuum S/N decreases, at some level the small-scale behaviour of the C-function becomes dominated by the noise so that the position of the "highest peak" may no longer be related to the actual Doppler shift, possibly resulting in a very large measurement error that cannot be predicted by the above estimators. This is illustrated in the bottom row of Fig. 3 where the extremum values for the early-type star obviously do not indicate the correct Doppler shift.

This phenomenon is well-known in signal analysis where it is sometimes referred to as "ambiguity". If the signal has a simple shape such as a radar pulse, the S/N level that marks the onset of ambiguity can be predicted theoretically (Woodward & Davies 1950) but unfortunately a stellar spectrum containing lines of different widths and strengths, many of them blended, allows no such prediction; this is illustrated by a comparison of the columns 1 and 2 or 3 and 4 in Fig. 3: for the early-type spectrum we find significant ambiguity already at $S/N = 7$, whereas for the late-type object this would occur at S/N=3 and below[6].

Notice in particular the difference in the C-function profiles between the "cool" object and the "hot" one: in the former the large-scale behaviour is dominated by the CaII lines while the relatively narrow central part is produced by the numerous metal lines in the spectrum; the early-type spectrum, on the other hand, contains only a small number of narrow lines which produce a weak peak superposed on a very broad base dominated by the Paschen lines: this weak structure may yet allow a fairly accurate measurement if the S/N is sufficiently high, but it is easily obliterated otherwise.

Notice also that the C-function curvature at a typical "pure noise" extremum is not very different from the one at a proper peak; this partly explains why the random error is often badly underestimated in the case of ambiguity.

5.3. The Monte-Carlo error bar

In spite of the best possible efforts, one may be faced with random error sources (such as ambiguity) whose effect cannot be described by the estimators above. A possible way of studying such errors is to use the following Monte-Carlo approach. For a given instrument and wavelength sampling one considers a wide variety of spectral-type, $v \sin i$, and S/N combinations; for each of these one simulates a large number ($N_{\rm mc}$) of observations by adding noise to a spectrum with known source velocity; measuring the Doppler shift of these simulated spectra one obtains a sample of $N_{\rm mc}$ simulated errors:

$$e_i = v_{i,\rm meas.} - v_{i,\rm known} \ . \tag{18}$$

The dispersion of these differences constitutes an alternative estimate of the random error, which we shall henceforth refer to as the Monte-Carlo error bar $\sigma_{\rm mc}$. A robust estimate of this dispersion is obtained as the sample semi-interquantile range (Höhle & Höhle 2009)

$$\sigma_{\rm mc} = \frac{1}{2}\left(\hat{Q}_e(0.8413) - \hat{Q}_e(0.1587)\right) \tag{19}$$

that corresponds to the standard deviation if the simulated errors follow a normal distribution.

The quantity $\sigma_{\rm mc}$ is not intended to replace an individual error estimate (except when the latter is obviously invalid) but it may provide useful additional information on the large-scale variation of random errors as a function of the atmospheric parameters of the observed source. Therefore, if one intends to observe a wide variety of objects with a given telescope-instrument combination, it is worthwhile again to create an auxiliary database from which an approximate value of $\sigma_{\rm mc}$ for any observed object can be obtained by interpolation. This can be done by calculating $\sigma_{\rm mc}$ for the nodes on a coarse grid such as the one described in Sect. 4.4, but with an additional dimension corresponding to the S/N level so that (unless additional APs need to be considered) it is five-dimensional.

6. Testing the methods

6.1. The bias test

Bias in a Doppler-shift measurement may originate from several sources. Most of these are best detected by performing a measurement on many different noiseless spectra for which (within the limits of numerical accuracy) the exact outcome is known, looking for deviations from the latter that are large enough to violate one's accuracy requirements. However, one cannot exclude the possibility of some mechanism interfering with the propagation of random errors and causing the final error on the measurement to be biased. To test whether a given measurement technique is liable to such a bias, one can use again the set of Monte-Carlo simulated errors of Eq. (18) and determine their median as a robust estimate of the bias.

The standard deviation of the median can be estimated as $\sigma_{|\rm M|} = \sigma_{\rm mc} 1.2533/\sqrt{N_{\rm mc}}$ (see e.g. Kendall & Stuart 1969, Example 10.7). Then, from a two-tailed test one may conclude, at the significance level α, that the measurement is biased if

$$\frac{|M|\sqrt{N_{\rm mc}}}{1.2533\sigma_{\rm mc}} > F^{-1}(1 - \frac{\alpha}{2}) \ , \tag{20}$$

where F is the standard normal cumulative distribution function.

6.2. The zscore test

In principle precision can be assessed by means of the individual estimators discussed in Sect. 5.1, but naturally one

[6] It is appropriate here to point out that the quantity 'continuum S/N' that we use as an indicator of apparent brightness, corresponds to the mean flux level (i.e. # photons per wavelength bin) at the continuum. Therefore it should not be interpreted as a direct indication of the precision of information obtained from the spectrum. A realistic S/N value in the latter sense (e.g. the mean ratio of the depth of line-cores to the amplitude of the noise) could be 2-10 times smaller.

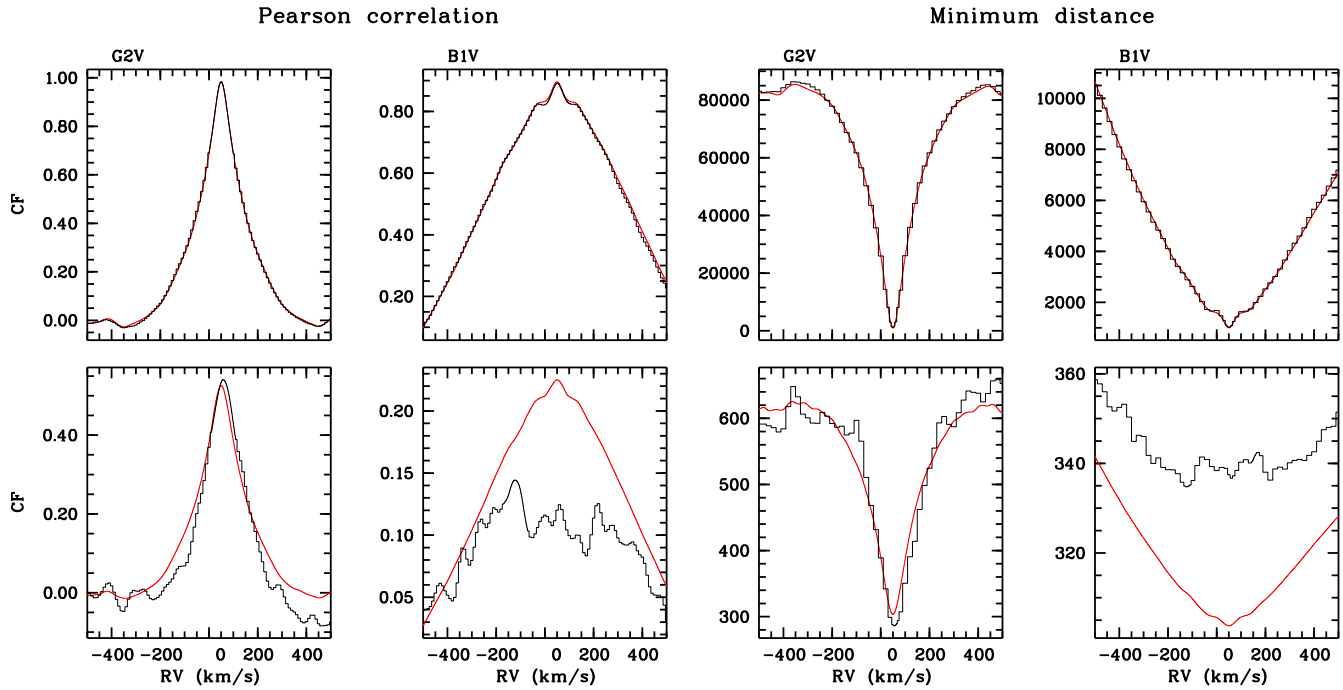

Fig. 3. C-functions on the wavelength range 847-874 nm, obtained with two of our methods for a late-type and an early-type object as indicated (synthetic spectra with only photon noise, no rotational broadening and "source" velocity = 50 km s^{-1}). Top row: spectrum grid step = 0.027 nm, continuum $S/N = 60$; bottom row: spectrum grid step = 0.088 nm, continuum $S/N = 7$. The red lines indicate the average over a large number of C-functions obtained from spectra differing only by their noise realization.

needs to verify to what extent these are reliable for the objects one intends to observe; this can be done as follows. We consider again a set of $N_{\rm mc}$ simulated errors (see Eq. (18)) and denote the random-error estimate accompanying each measurement by σ_i; then we define the quantity

$$z_i = \frac{e_i}{\sigma_i} \ , \tag{21}$$

which is termed the zscore of the measurement. If in fact the estimator correctly describes the true measurement error, the quantities z_i follow a standard normal distribution. This could be ascertained by means of e.g. the one-sample Kolmogorov-Smirnoff test, but a quicker (and, as we found, sufficiently reliable) approach consists in determining their dispersion σ_z and comparing this to its sampling distribution. Again we use the robust estimate

$$\sigma_z = \frac{1}{2}\left(\hat{Q}_z(0.8413) - \hat{Q}_z(0.1587)\right) \ . \tag{22}$$

The sampling distribution of σ_z for a standard normal variable is normal with mean=1 and standard deviation stdev$[\sigma_z] \approx 0.962/\sqrt{N_{\rm mc}}$ (see e.g. Kendall & Stuart 1969, Example 10.8 with $p_1 = 0.8413$ etc.). Then, from a two-tailed test one may conclude, at the significance level α, that the estimator is not valid if

$$\frac{|\sigma_z - 1|\sqrt{N_{\rm mc}}}{0.962} > F^{-1}(1 - \frac{\alpha}{2}) \ . \tag{23}$$

6.3. Comparing performances

Consider any two measurement methods that both satisfy the above quality requirements; then naturally the question arises whether one of them nevertheless induces somewhat larger errors than the other and, more to the point, under what circumstances this is the case. To investigate this we consider, again, sets of Monte-Carlo simulated errors as defined in Eq. (18), pairing the samples so that we always compare two measurements obtained with the same noise realization. Although we shall frequently use the terms "better" or "best" in the following discussions, this will never be meant in an absolute sense, firstly because we only consider the magnitude of the errors (disregarding e.g. the computational load of a given method) and secondly because the test results are likely to depend on resolution and sampling, on the noise level, the spectral type, and the rotational broadening of the objects. A different choice for these parameters could easily lead to a different conclusion.

Supposing one wishes to test separately for differences in bias and in dispersion, one could proceed as follows. In a set of Monte-Carlo simulated errors the bias appears as the mean of these errors; a classical test to detect a difference in the means of two normally distributed populations is based on Student's t-distribution and it has a version that is applicable to paired samples. However, for the dispersion the situation is different: we found no paired-sample counterpart for the classical F-test on the homogeneity of variances; therefore we propose to use again the t-test, but now on the absolute value of the errors after subtraction of the bias, as detailed below.

These are not the only possibilities of course. If one is interested e.g. in finding out whether there is a significant difference in the total error (bias + random part) one could repeat the second test without subtracting the bias. And if one is worried about the fact that a difference of absolute values of errors may not be normally distributed, one could replace the t-test by a non-parametric one such as the Sign test (see e.g. Siegel 1956).

6.3.1. Bias

Let us name the methods, arbitrarily, "1" and "2"; we note $e_{1,i}$ and $e_{2,i}$ for the respective error incurred with each of the $N_{\rm mc}$ simulated spectra and we define the differences

$$d_i = e_{1,i} - e_{2,i} \ . \tag{24}$$

Next we compute the average $\overline{d}$ and the standard deviation $\sigma_{\rm d}$. If both methods perform equally well as far as the bias is concerned, the population mean of d equals zero; furthermore, provided that both $e_{1,i}$ and $e_{2,i}$ follow a Gaussian distribution with the same dispersion, the statistic $t_{\rm d} = \overline{d}/(\sigma_{\rm d}/\sqrt{N_{\rm mc}})$ is distributed according to a Student's t-distribution with $N_{\rm mc} - 1$ degrees of freedom (see e.g. Alder & Roessler 1972). If $N_{\rm mc}$ is sufficiently large (e.g. ≥ 500) the t-distribution may be replaced by the standard normal. Then, from a two-tailed test one concludes, at the significance level α, that this null-hypothesis must be rejected if

$$\frac{\overline{d}\sqrt{N_{\rm mc}}}{\sigma_{\rm d}} > F^{-1}(1-\frac{\alpha}{2}) \ . \tag{25}$$

Subsequently a comparison between $|\overline{e_1}|$ and $|\overline{e_2}|$ will indicate which of the two methods causes the smallest bias. Note that, if $\overline{e_1}$ and $\overline{e_2}$ have different signs, the statistic may indicate a significant difference without either of the methods being truly "better" than the other. Also, even if they have the same sign and a statistically significant difference, it is possible that this difference is not physically relevant because the largest bias is still sufficiently small to be harmless for the required accuracy.

6.3.2. Dispersion

For the present purpose we use the mean absolute deviation to characterize the dispersion. Let $e'_{1,i} = e_{1,i} - \overline{e_1}$, $e'_{2,i} = e_{2,i} - \overline{e_2}$, $d'_i = |e'_{1,i}| - |e'_{2,i}|$ and compute the average $\overline{d'}$ and the standard deviation $\sigma_{\rm d'}$. If both methods produce measurements with the same dispersion, we expect that the population mean of d' equals zero. Furthermore we assume that the differences d'_i are normally distributed so that the statistic $t_{\rm d'} = \overline{d'}/(\sigma_{\rm d'}/\sqrt{N_{\rm mc}})$ follows a t-distribution with $N_{\rm mc} - 1$ degrees of freedom. Then, as for the bias-difference, we can conclude that the difference $\overline{|e'_1|} - \overline{|e'_2|}$ is significant at the level α if

$$\frac{\overline{d'}\sqrt{N_{\rm mc}}}{\sigma_{\rm d'}} > F^{-1}(1-\frac{\alpha}{2}) \ , \tag{26}$$

and in that case, see which method causes the smallest mean absolute deviation.

6.3.3. Application

To illustrate the use of these tests, we apply them to a set of spectra corresponding to different values of effective temperature and continuum S/N ratio; otherwise all spectra have $\log g = 4.0$, $v \sin i = 0.0\,{\rm km\,s^{-1}}$, and solar abundances. The spectra simulate observations from the Gaia-RVS instrument: they cover the wavelength range 847–874 nm with an instrumental resolution of 11,500; the wavelength sampling-step is 0.027 nm for the brighter spectra (continuum $S/N > 25$) and 0.088 nm otherwise. For each combination of $T_{\rm eff}$ and S/N, a sample of $N_{\rm mc} = 1000$ spectra was generated with different noise realizations.

In Fig. 4 we compare the performance of the methods two by two. In order to be able also to judge the physical relevance of the differences, we scaled the size of the symbols with the difference in bias magnitude

$$s_{\rm d} = ||\overline{e_1}| - |\overline{e_2}|| \tag{27}$$

or with the difference in standard deviation

$$s_{\rm d'} = |{\rm stdev}(e_1) - {\rm stdev}(e_2)| \ , \tag{28}$$

both expressed in ${\rm km\,s^{-1}}$. Note that larger differences are not necessarily also statistically more significant.

For these specific data, there are clear areas in the ($T_{\rm eff}$, S/N) diagram where one technique provides a significantly sharper dispersion than the other. Differences between the techniques also exist for the bias, but these seem to be less consistently grouped into areas of the diagram and they are, on the whole, statistically less significant. We caution again that the conclusions derived here apply to the specific instrument for which these simulated data were calculated. Other instrument configurations (wavelength range, resolution, ...) could easily lead to different conclusions.

6.4. Global prognostics for a survey

The various tests described above do not involve observed data so they can be performed independently of the actual measurements, though for a given telescope-instrument combination. Thus, if one expects to observe a wide variety of spectra it is best to perform them as early as possible and to store their results in an auxiliary database (ADB) such as those already mentioned in Sects. 4.4 and 5.3.

In a multi-method approach, all measurement algorithms are naturally implemented on a common platform where most of the peripheral operations (e.g. data preparation and template construction) are performed by separate modules communicating with the several measurement modules through appropriate interfaces. Such an architecture favours the implementation of a dedicated "tests" module that

- defines a common (coarse) AP test grid[7];
- generates $N_{\rm mc}$ simulated spectra for each node;
- calculates the Monte-Carlo error bar $\sigma_{\rm mc}$ and the left hand side of Eqs. (20,23) with each of the measurement methods;
- for each couple of methods, calculates the left hand side of Eqs. (25,26) and the size differences (Eqs. (27,28));
- stores all results in the ADB.

It is convenient to calculate the coefficients of the local model for the TME (Eq. (15)) on the same grid (but without the (S/N) dimension) so they can be stored in the same ADB.

The "tests" module should also offer appropriate tools for querying the ADB and for plotting, so that graphs such as Fig. 2 or Fig. 4 can be produced easily. Thus from this database one can obtain an overall picture of the expected quality of the Doppler-shift measurements. This in turn may provide useful feedback to further improve the implementation of the methods and/or to modify the observational set-up or even to avoid observations that are likely to be profitless.

[7] e.g. as described in Sect. 5.3

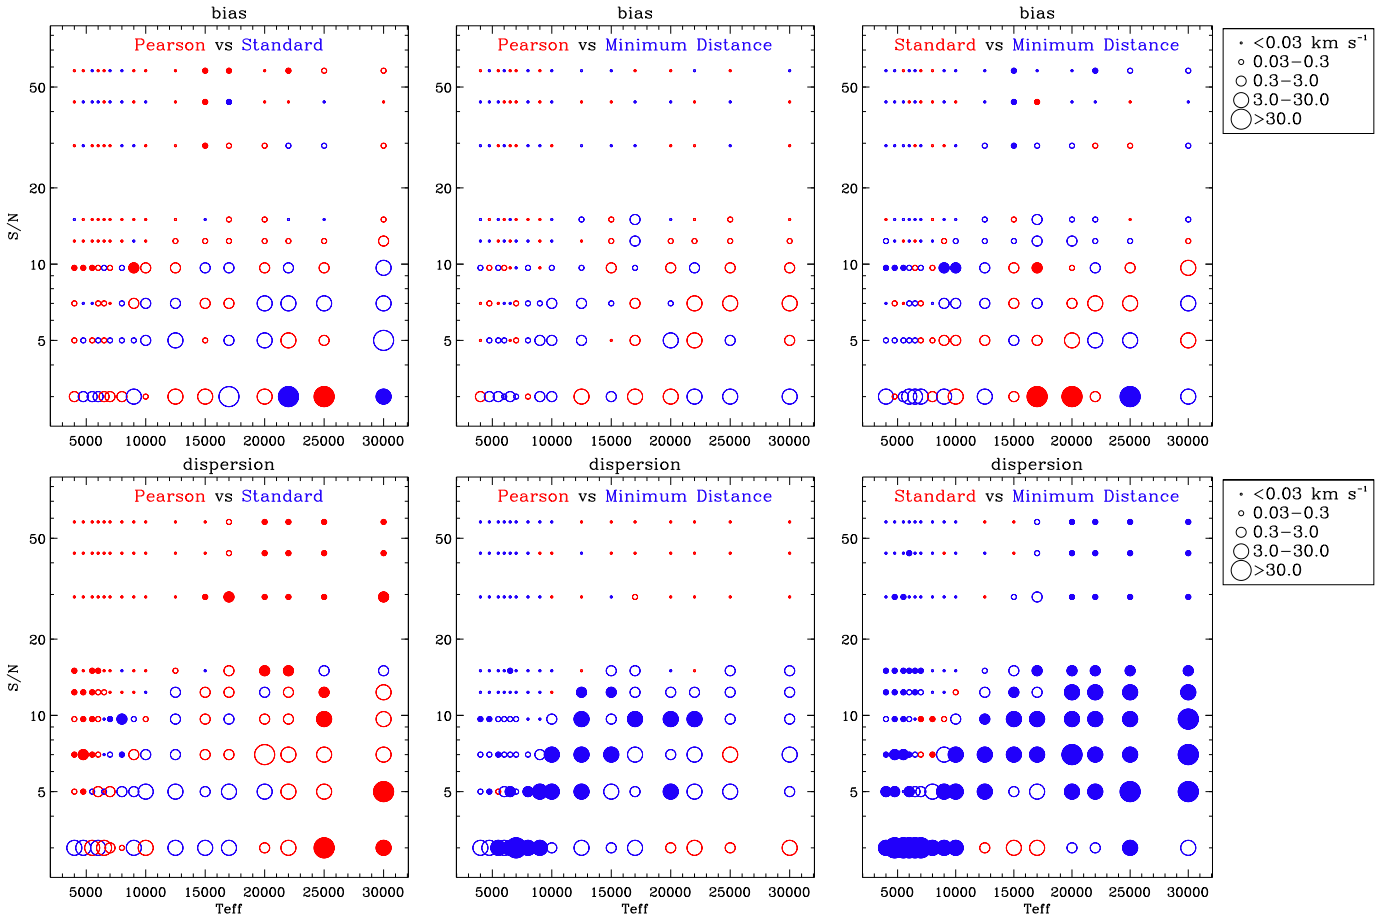

Fig. 4. Pairwise comparison of the bias (top panels) and the dispersion (bottom panels) between three different methods as indicated in each panel. The results are plotted as a function of effective temperature (x-axis) and signal-to-noise ratio (y-axis). The colours are identified with the method names in each panel; each point is given the colour of the method with the smallest bias (top row) or the smallest dispersion (bottom row) respectively. Filled symbols indicate that the difference is statistically significant at the 0.2% level in a two-sided test.

7. Combining the measurements

The user of any measurement software naturally wants a single answer within some confidence region. Assume now that we have an auxiliary database as described above and that a (common) significance level α for the various hypothesis tests has been chosen; then, for each observed spectrum we can proceed as follows:

- determine which node on the grid best matches the observation and retrieve the corresponding data to be used in the following steps;
- select the methods that pass the bias test; the measurement produced by a method that fails this test should be stored but not used in a combination;
- if any of the remaining methods fails the zscore test, its internal error estimate must be replaced by its Monte-Carlo error bar;
- considering the TME as a random effect (see Sect. 4.1), estimate its magnitude[8] using the local-model coefficients for each remaining method and add it quadratically to the random error;
- obtain the combined measurement as the median or the error-weighted average of the remaining individual measurements;
- estimate the combined error using standard error propagation;

The above procedure is not very refined, but at least its automation is feasible and it should provide more reliable answers than a simple average of all measurements. Obviously its success will hinge on the richness and the quality of the auxiliary database.

Nevertheless one should remain aware of the fact that there may be situations (e.g. when the various measurements all disagree within the limits of their individual uncertainty, as is not inconceivable with faint objects) where one could hardly have confidence in the result of any automated combination, however sophisticated be its scheme. Therefore it is best to keep on record not only the combined value and its error estimate, but also the result produced by each method separately. In this way the user can judge whether it is necessary to have a closer look at the individual results and perhaps combine these in another way or to go back even further and handle the measurement manually for the case at hand.

[8] a safe option is to take the maximal magnitude of the TME over the unit cell (centred on the object's APs) of the fine grid defined in Sect. 4.2.1

8. Conclusions

Among the numerous methods for Doppler-shift measurement available in the literature, we select three that do not impose strong requirements on the observed data, regarding resolution and S/N. For one of these we provide two variants that one may use lacking a realistic estimate of the error on each flux sample. The theoretical foundation of the methods is discussed to provide a better understanding of any differences in the results they yield.

With each method we select an appropriate estimator for the internal error on the measurements and we propose a Monte-Carlo based global estimator, to be used e.g. with very faint objects where the internal estimator becomes unreliable. We also discuss in some detail how the so-called template mismatch error can be dealt with in a systematic way. Several tests allowing to judge the performance of the methods are described.

In our multi-method approach we propose a combination of the individual Doppler-shift measurements, using an auxiliary database that roughly predicts the measurement-quality one can expect for any type of object to be observed. If such a database can be established already at an early stage in a survey project, it may moreover provide useful information to optimize the instrumental set-up and the observational strategy or even the instrumental design.

Acknowledgements. We acknowledge funding by the Belgian Federal Science Policy Office through ESA PRODEX programme "Binaries, extreme stars and solar system objects", under contract nrs. C90289 and C90290. Work done in GEPI has been partly funded by CNES through programme "CNES-Gaia" under contract nr. 92532/o/. The authors would like to thank the Gaia Data Processing and Analysis Consortium, and in particular the colleagues in Coordination Unit 6 ("Spectroscopic Processing"), for numerous fruitful discussions.

Appendix A: Template sampling in $\ln\lambda$ space

If the wavelength grid has a constant step Δ in $\ln\lambda$, the central wavelength of the bins can be written as

$$\lambda_n = \lambda_0\, e^{n\Delta} \ . \tag{A.1}$$

Since the sampling Eq. (1) is contiguous we have $\lambda_n + \Delta_n = \lambda_{n+1} - \Delta_{n+1}$ so that, noting $\mu_n = \lambda_n - \Delta_n$ we can write $\lambda_n + \Delta_n = \mu_{n+1}$ and thus $\lambda_n = \frac{\mu_n + \mu_{n+1}}{2}$; then it is easily seen that

$$\mu_n = \mu_0 e^{n\Delta} \quad , \quad \mu_0 = \frac{\lambda_0\, e^{-\frac{\Delta}{2}}}{\cosh(\frac{\Delta}{2})} \ . \tag{A.2}$$

Consider now in particular the discrete set of velocities satisfying

$$1 + \frac{v_m}{c} = e^{m\Delta} \quad , \quad m \in \mathbb{N} \ ; \tag{A.3}$$

then, provided the approximation Eq. (5) is valid:

$$\begin{aligned} t_n(v_m) &= \int_{\mu_n e^{-m\Delta}}^{\mu_{n+1} e^{-m\Delta}} t(\lambda;0)\mathrm{d}\lambda = \int_{\mu_{n-m}}^{\mu_{n-m+1}} t(\lambda;0)\mathrm{d}\lambda \\ &= t_{n-m}(0) \ . \end{aligned} \tag{A.4}$$

Thus we obtain the familiar picture of a rigid shift of the spectrum as a result of the motion of the source. This implies in particular that for different velocities satisfying Eq. (A.3), even the simplified sampling in Eq. (5) does not have to be recalculated for each velocity.

If other velocities need to be considered, it may be useful to rewrite Eq. (1) in a form that closely resembles Eq. (A.4) by defining a new variable s:

$$1 + \frac{v}{c} = e^{s\Delta} \tag{A.5}$$

so that $t_n(v)$ can be written as a function of $n - s$:

$$t_n(v) = h(n-s) = \int_{\mu_0 e^{(n-s)\Delta}}^{\mu_0 e^{(n+1-s)\Delta}} t(\lambda;0)\mathrm{d}\lambda \ . \tag{A.6}$$

Thus Eq. (A.5) defines the Doppler shift as a real quantity, rather than an integer as in Eq. (A.3)

Appendix B: Other flux distributions

If a source is sufficiently bright one can assume that the PDF is Gaussian and estimate its dispersion as $\sigma_n^2 = at_n(v) + \sigma_{\mathrm{ro}}^2$ where σ_{ro} represents e.g. "read-out noise" or any constant additive noise source. Note that, unlike σ_n in Eq. (10), this error model depends on the template velocity. Thus, instead of Eq. (10) we now have to minimize

$$C_{\mathrm{mdp}}(a,v) = \sum_{n=1}^{N} \left\{ \frac{(f_n - at_n(v))^2}{at_n(v) + \sigma_{\mathrm{ro}}^2} + \ln(at_n(v) + \sigma_{\mathrm{ro}}^2) \right\} \ . \tag{B.1}$$

An exact 1-dimensional version of this function has not been found. However, if the scale factor a is known independently we can put $a = 1$ and then, neglecting the logarithmic term and assuming there is only photon noise, Eq. (B.1) becomes the so-called S-statistic (see e.g. Cash 1979; Lampton et al. 1976, and references therein). It should be noted that the error model assumed above cannot be used down to arbitrarily low flux levels. In such cases one would have to build a somewhat more sophisticated PDF that properly accounts for the combination of photon noise and an additional Gaussian noise source.

If all noise sources except photon counting are negligible, the flux-sample values are Poisson-distributed, i.e.

$$\mathrm{pd}(f_n|a,v) = \frac{(at_n(v))^{f_n} e^{-at_n(v)}}{f_n!} \ . \tag{B.2}$$

In that case, maximizing the likelihood is equivalent to minimizing

$$C_{\mathrm{mdc}}(a,v) = 2\sum_{n=1}^{N} \{at_n(v) - f_n \ln(at_n(v))\} \ , \tag{B.3}$$

where the factor 2 is introduced to allow for the application of Eq. (17). At the extremum we now have

$$\hat{a} = \frac{\sum f_n}{\sum t_n(\hat{v})} \ , \tag{B.4}$$

which can be substituted in Eq. (B.3) to obtain a 1-dimensional C-function. On the other hand, if again the parameter a can be omitted, the right hand side of Eq. (B.3) equals the so-called C-statistic (Cash 1979). The C-function from Eq. (B.3) is especially suitable when the observations

deliver a pure photon count of a very weak source, without any other noise sources to be considered.

From a performance-comparison as in Fig. 4 we found that $C_{\rm md}$ is preferable but that both $C_{\rm mdp}$ and $C_{\rm mdc}$ are viable alternatives if, for some reason, $C_{\rm md}$ cannot be used.

To the best of our knowledge, neither of the functions $C_{\rm mdp}$ and $C_{\rm mdc}$ has been used before in radial velocity measurements. We note, however, that the problem of Doppler-shift measurement is mathematically very similar to the detection of periodicity in time-series analysis, for which the MLP method (Zane et al. 2002) does use the C-statistic.

Appendix C: Upper bound for the MME

Appendix C.1: Symmetric functions

We note x for the independent variable, y for the corresponding function value, s for the sampling step size, (x_0, y_0) for the data point with the highest/lowest sample value within the range of x and $(x_0 \pm s, y_\pm)$ for the nearest adjacent samples so that

$$(y_0 - y_+)(y_0 - y_-) \geq 0 \ . \tag{C.1}$$

Furthermore we note $x_{\rm c}$ for the true position of the extremum and $x_{\rm p}$ for the peak-position estimated using a parabola as interpolating function. Then

$$x_{\rm p} = x_0 - \left(\frac{y_+ - y_-}{y_+ + y_- - 2y_0} \right) \frac{s}{2} \ . \tag{C.2}$$

Even if the function is intrinsically symmetric this estimate has a modest MME related to the discrete nature of the sampling, i.e. there is a difference

$$\delta = x_{\rm p} - x_{\rm c} \tag{C.3}$$

that depends on the actual shape of the function. It is easily seen that $\delta = 0$ in the special cases $y_+ = y_-$ and $y_0 = y_\pm$; otherwise, assuming that the function has no structure on a scale less than s and that there is no noise, one finds that $|\delta| < \frac{s}{2}$. Of course the error will be much less than this if the function closely resembles a parabola.

Anyway, it follows that $|\delta|$ could be made arbitrarily small by reducing the step size, as is possible with the velocity-space methods in Sect. 2.3. However, in the standard method the step size s cannot be reduced at will so there we reduce the basic MME by combining the above three-point fit with a four-point fit following David & Verschueren (1995).

Appendix C.2: Weakly asymmetric functions

We assume that s is sufficiently small to ensure that, within the interval $[x_0 - 2s, x_0 + 2s]$, the function to be centroided is well approximated by a third-degree polynomial:

$$y(x) = \alpha + \beta(x - x_{\rm c})^2 + \gamma(x - x_{\rm c})^3 \tag{C.4}$$

where the parameter β can be said to control the contrast of the extremum value with respect to the adjacent points and γ to control the asymmetry[9]; if Eq. (C.4) represents a C-function, $1/\beta$ reflects the width of the narrowest lines present in the spectra (see also Verschueren & David 1999).

Then, with the notations introduced above, choosing x_0 according to (C.1), after some algebra one finds that

$$|\delta| < \frac{\left|\frac{\gamma s}{\beta}\right| s}{1 - 3\left|\frac{\gamma s}{\beta}\right|} \quad \text{if} \quad s < |\frac{\beta}{3\gamma}| \ . \tag{C.5}$$

[9] Of course this interpretation is meaningful only if $|\gamma s| << |\beta|$; as an example, for the red lines in the top row of Fig. 3, one has $5.10^{-4} < |\frac{\gamma s}{\beta}| < 10^{-2}$.

References

Alder, H. L. & Roessler, E. B. 1972, Introduction to probability and statistics (San Francisco, CA: W. H. Freeman and Co.)
Allende Prieto, C. 2007, AJ, 134, 1843
Andersen, J. & Nordstrom, B. 1983, A&A, 122, 23
Baranne, A., Mayor, M., & Poncet, J. L. 1979, Vistas in Astronomy, 23, 279
Baranne, A., Queloz, D., Mayor, M., et al. 1996, A&AS, 119, 373
Behr, B. B., Hajian, A. R., Cenko, A. T., et al. 2009, ApJ, 705, 543
Bendat, J. S. & Piersol, A. G. 1971, Random Data: Analysis and Measurement Procedures (New York: Wiley-Interscience)
Bouchy, F., Pepe, F., & Queloz, D. 2001, A&A, 374, 733
Brown, T. M. 1990, in A. S. P. Conf. Ser., Vol. 8, CCDs in astronomy, ed. G. H. Jacoby, 335–344
Butler, R. P., Marcy, G. W., Williams, E., et al. 1996, PASP, 108, 500
Cahoy, K., Fischer, D., Spronck, J., & DeMille, D. 2010, in Proceedings of SPIE-The International Society for Optical Engineering, Vol. 7739, Modern Technologies in Space- and Ground-Based Telescopes and Instrumentation, ed. E. L. D. Atad-Ettedgui
Cash, W. 1979, ApJ, 228, 939
Chelli, A. 2000, A&A, 358, L59
Chen, C. H., Mamajek, E. E., Bitner, M. A., et al. 2011, ApJ, 738, 122
Connes, P. 1985, Ap&SS, 110, 211
Cornu, A. 1890, Sur la méthode Doppler-Fizeau permettant la détermination, par l'analyse spectrale, de la vitesse des astres dans la direction du rayon visuel (Paris: [Impr. Gauthier-Villars et fils, 1890])
Da Costa, G. S., Freeman, K. C., Kalnajs, A. J., Rodgers, A. W., & Stapinski, T. E. 1977, AJ, 82, 810
David, M. & Verschueren, W. 1995, A&AS, 111, 183
De Loore, C., Monderen, P., & Rousseeuw, P. 1987, A&A, 178, 307
Deslandres, H. 1898, Astronomische Nachrichten, 148, 23
Díaz, C. G., González, J. F., Levato, H., & Grosso, M. 2011, A&A, 531, A143
Dravins, D., Lindegren, L., & Madsen, S. 1999, A&A, 348, 1040
Fekel, F. C. 1999, in A. S. P. Conf. Ser., Vol. 185, IAU Colloq. 170: Precise Stellar Radial Velocities, ed. J. B. Hearnshaw & C. D. Scarfe, 378–+
Fellgett, P. 1955, Optica Acta, 2, 9
Furenlid, I. & Furenlid, L. 1990, PASP, 102, 592
Gray, D. F. 2005, The Observation and Analysis of Stellar Photospheres (Cambridge University Press)
Griffin, R. E. M., David, M., & Verschueren, W. 2000, A&AS, 147, 299
Griffin, R. F. 1967, ApJ, 148, 465
Griffin, R. F. & Gunn, J. E. 1974, ApJ, 191, 545
Groote, D. & Hunger, K. 1977, A&A, 56, 129
Gullberg, D. & Lindegren, L. 2002, A&A, 390, 383
Hill, G. 1982, Publications of the Dominion Astrophysical Observatory Victoria, 16, 59
Hill, G. 1993, in A. S. P. Conf. Ser., Vol. 38, New Frontiers in Binary Star Research, ed. K.-C. Leung & I.-S. Nha, 127–+
Höhle, J. & Höhle, M. 2009, {ISPRS} Journal of Photogrammetry and Remote Sensing, 64, 398
Huggins, W. 1868, Royal Society of London Philosophical Transactions Series I, 158, 529
Jenkins, G. M. & Watts, D. G. 1969, Spectral analysis and its applications (London: Holden-Day)
Kambe, E., Ando, H., Sato, B., et al. 2008, Publications of the Astronomical Society of Japan, 60, 45
Katz, D., Munari, U., Cropper, M., et al. 2004, MNRAS, 354, 1223
Kendall, M. & Stuart, A. 1969, The advanced theory of statistics. Vol.1: Distribution theory (London: Charles Griffin & Co. Ltd., 3rd ed.)
Klinkerfues, E. F. W. 1866, Astronomische Nachrichten, 66, 337

Kurtz, M. J. & Mink, D. J. 1998, PASP, 110, 934
Lampton, M., Margon, B., & Bowyer, S. 1976, ApJ, 208, 177
Li, C.-H., Benedick, A. J., Fendel, P., et al. 2008, Nature, 452, 610
Lindegren, L. & Dravins, D. 2003, A&A, 401, 1185
Lindegren, L., Madsen, S., & Dravins, D. 2000, A&A, 356, 1119
Lovis, C. & Pepe, F. 2007, A&A, 468, 1115
Mahadevan, S. & Ge, J. 2009, ApJ, 692, 1590
Martin, B. R. 1971, Statistics for physicists (London: Academic Press)
Mazeh, T. & Zucker, S. 1992, in A. S. P. Conf. Ser., Vol. 32, IAU Colloq. 135: Complementary Approaches to Double and Multiple Star Research, ed. H. A. McAlister & W. I. Hartkopf, 164–+
Mazeh, T., Zucker, S., Goldberg, D., et al. 1995, ApJ, 449, 909
Murdoch, K. & Hearnshaw, J. B. 1991, Ap&SS, 186, 137
Murphy, M. T., Udem, T., Holzwarth, R., et al. 2007, MNRAS, 380, 839
North, R. C., Marsh, T. R., Kolb, U., Dhillon, V. S., & Moran, C. K. J. 2002, MNRAS, 337, 1215
Penrose, R. 1955, Mathematical Proceedings of the Cambridge Philosophical Society, 51, 406
Penrose, R. 1956, Mathematical Proceedings of the Cambridge Philosophical Society, 52, 17
Press, W. H., Flannery, B. P., & Teukolsky, S. A. 1986, Numerical recipes. The art of scientific computing
Ramm, D. J. 2004, PhD thesis, University of Canterbury
Reiners, A., Bean, J. L., Huber, K. F., et al. 2010, ApJ, 710, 432
Royer, F. 1999, PhD thesis, Observatoire de Paris
Royer, F., Gerbaldi, M., Faraggiana, R., & Gómez, A. E. 2002, A&A, 381, 105
Sargent, W. L. W., Schechter, P. L., Boksenberg, A., & Shortridge, K. 1977, ApJ, 212, 326
Seifahrt, A. & Kaeufl, H. U. 2008, A&A, 491, 929
Siegel, S. 1956, Non-parametric statistics for the behavioral sciences (New York: McGraw-Hill)
Simkin, S. M. 1974, A&A, 31, 129
Sohncke, L. 1867, Astronomische Nachrichten, 69, 209
Steinmetz, M., Siebert, A., Zwitter, T., & RAVE Collaboration. 2009, in IAU Symposium, Vol. 254, IAU Symposium, ed. J. Andersen, J. Bland-Hawthorn, & B. Nordström, 453–460
Steinmetz, M., Zwitter, T., Siebert, A., et al. 2006, AJ, 132, 1645
Steinmetz, T., Wilken, T., Araujo-Hauck, C., et al. 2008, Science, 321, 1335
Tonry, J. & Davis, M. 1979, AJ, 84, 1511
Torres, G., Latham, D. W., & Stefanik, R. P. 2007, ApJ, 662, 602
Valdivielso, L., Esparza, P., Martin, E. L., Maukonen, D., & Peale, R. E. 2010, ApJ, 715, 1366
van Eyken, J. C., Ge, J., & Mahadevan, S. 2010, ApJS, 189, 156
Verschueren, W. 1991, PhD thesis, Vrije Universiteit Brussel
Verschueren, W. & David, M. 1999, A&AS, 136, 591
Verschueren, W., David, M., & Griffin, R. E. M. 1999, A&AS, 140, 107
Weiss, W. W., Jenkner, H., & Wood, H. J. 1978, A&A, 63, 247
Wilkinson, M. I., Vallenari, A., Turon, C., et al. 2005, MNRAS, 359, 1306
Woodward, P. M. 1953, Probability and Information Theory, with Applications to Radar (London: Pergamon Press)
Woodward, P. M. & Davies, I. L. 1950, Philosophical Magazine, 41, 1001
Zane, S., Haberl, F., Cropper, M., et al. 2002, MNRAS, 334, 345
Zucker, S. 2003, MNRAS, 342, 1291
Zucker, S. & Mazeh, T. 1994, ApJ, 420, 806
Zverko, J., Ziznovsky, J., Mikulasek, Z., & Iliev, I. K. 2007, Contributions of the Astronomical Observatory Skalnate Pleso, 37, 49
Zwitter, T., Munari, U., & Siebert, A. 2005, in ESA Special Publication, Vol. 576, The Three-Dimensional Universe with Gaia, ed. C. Turon, K. S. O'Flaherty, & M. A. C. Perryman, 623–+